Super-Intelligence and the virtues of a "Hive Mind"

Louis Rosenberg / Op Ed

Posted on: February 10. 2016 / **Last Modified:** February 10, 2016

We humans pride ourselves on being rational thinkers with an inherent sense of morality that guides our actions towards the greater good. These virtues hold true across all levels of society and yet collectively. on a global scale, we often make self-destructive decisions. I'm talking about the kind of decisions that lead to war, pollution, poverty. inequality, and in recent years, climate change.

This begs the question, how can immoral decisions emerge from a society comprised overwhelmingly of moral individuals? Philosophers have been pondering this for ages. Nietzsche lamented – *"Madness is rare in individuals – but in groups, political parties, nations, and eras it's the rule."* Renowned American theologian, Reinhold Niebuhr was even more blunt, expressing –*"the group is more arrogant, hypocritical, self-centered, and more ruthless in the pursuit of its ends than the individual."* So, what is it about human groups that cause us to behave so differently together than we would behave alone?

Social scientists often cite the "Tragedy of the Commons" problem when pondering group morality. First postulated by the Victorian economist William Foster Lloyd in 1833, the premise is that individuals, who act both morally and rationally on a local scale, are prone to producing immoral results on a group scale. He pointed to herdsman running cattle on open pastures. As an individual rancher, it's entirely rational and moral to maximize the size of your herd. But, if *all* herdsman follow this individual morality, the shared pasture

gets overrun and is ruined for all. Thus individual morality is not always aligned with the common good. In fact, it may be misaligned more often than not.

Ecologist Garett Hardin brought this to modern relevance in a 1968 when he linked this to population growth in the journal Science. He pointed out that on a local level, it's a basic Human Right for parents to decide the number of children to have. And for much of the world's inhabitants, a large brood is fully rational, optimizing survival of the family. On a global level, however, if all families behave under that same local morality, overpopulation will likely result, putting most families in danger.

So, how do we better handle social dilemmas in which the short-term interests of individuals are at odds with the long-term interests of the group? To date, the most successful path has been the use of democratic governance in which groups make decisions collectively, through direct or representative polling of the population. The presumption is that by revealing the consequences of their collective actions to the full group, democratic decisions will emerge that support the common good. The problem is, our current methods for polling groups often fall victim to the "Tragedy of the Commons" pitfall.

A clever example of this was recently performed at University of Maryland by Dylan Selterman. He posed an extra-credit challenge to his Social Psychology class, allowing each of his students to indicate by secret ballot how many points of extra credit they wanted on their exam – 2 points or 6 points. The only twist was that if more than 10% of the class asked for 6 points, nobody would get any bonus. Clearly, it was in the best interest for everyone in the class to individually ask for 2 points, but that's not what happened. Far too many students asked for 6 points and nobody received extra credit.

So, are we humans doomed to make self-destructive global decisions because of something flawed within us? Or is the Tragedy of the Commons problem a consequence not of our nature, but of our methods for group

decision making? An optimist, my view is that our tendency for self-destructive decisions is not because of a fundamental human flaw, but because our modern decision-making process is broken – the way we mediate opposing interests, weigh competing alternatives, and converge on final outcomes. The fact is, our current methods are highly influenced by special interests, the more extreme the position the more attention given, thereby producing solutions that not optimal for the common good.

And the problem is getting worse, for we've become a "poll obsessed" society, overusing a crude tool meant for quantifying groups, while forgetting that polls do little to encourage consensus or help groups reach smart decisions that support the common good. Much the opposite, polls usually are polarizing, highlighting the differences in a population, while encouraging special-interests to entrench. This is why rational and moral individuals are often unable to agree on solutions that are best for the population at large, even in a democracy that aims to achieve this. Instead, we either stagnate with no decision being reached, or we polarize, entrenching around positions that go against the group's long-term self-interest.

So, is there a way to encourage rational and moral group decisions? We could look to Mother Nature for guidance. Countless species have evolved methods for quickly reaching group decisions based on input from large populations of diverse individuals. From schools of fish and flocks of birds, to colonies of ants and swarms of bees, nature achieves this feat, not by taking votes or polls, but by enabling groups to form *real-time dynamics systems* that negotiate in synchrony and converge on optimal outcomes.

Biologists call the phenomenon Swarm Intelligence. It's the way nature has learned to tap into the diverse knowledge, intuition, experiences, and instincts of groups and produce decisions that are better for the common good than could be produced by any single individual. In fact, research has shown that swarming amplifies the intelligence of the species, resulting in "super-organisms" that can solve problems and make decisions that are beyond the

capacity of the individual members.

Fig 1 "Swarm of Honeybees"

The most deeply researched swarms in nature are those formed by honeybees. As studied in detail by Tomas Seeley at Cornell University. the decision-making process of honeybee swarms has been shown to remarkably similar to that performed by neurological brains. Both employ large populations of simple units (i.e., bees and neurons) that work in parallel to integrate noisy evidence, weigh competing alternatives, and converge on decisions in synchrony. In this sense, Honeybees and other organisms that swarm are able make decisions by forming a *"brain of brains"* that is far more intelligent than any individual contributor.

For example, honeybees face a life-or-death decision when selecting a location for a new colony. After searching a 30 square mile area, scout bees bring information about dozens of potential sites back to the swarm for consideration. Each site is assessed across many competing criteria, including –

safety from predators, insulation for winter, ventilation for summer, and storage capacity for honey. Using body vibrations known as a "waggle dance", the scout bees form a real-time s arm where they express preferences for various sites based on the many quality factors. Through dynamic negotiation among the competing signals, a decision is reached. The amazing thing is that hone bees as studied by Seeley converge on the optimal decision 80% of the time. They don't get mired in stagnation and indecision. The don t ge hijacked by special interest groups. Instead, the pool their diverse knowledge and preferences and through the natural process of swarming, efficiently reach an optimized decision that is best for the survival of the group as a whole.

So what is so terrible about a "*hive mind?*" I suspect the negative connotations are primarily a consequence of deep misconceptions about bees.

any people assume that bees are "drones" that take blind direction from an all-powerful queen. This is simply incorrect. The queen has no influence on colon decisions. Instead honeybees make decisions by concerning a swarm of 300 to 500 of their most experienced scouts who negotiate in real-time weighing the alternatives in a democratic and thoughtful manner. In many ways their process is less ' drone-like" than elections held by us humans wherein most participants reflexively vote along party lines the decisions being made by a small percentage of independents in the middle. The fact is, bees negotiate and compromise while e polarize and entrench. I would argue that nature s "hive minds" are more enlightened than we humans appreciate.

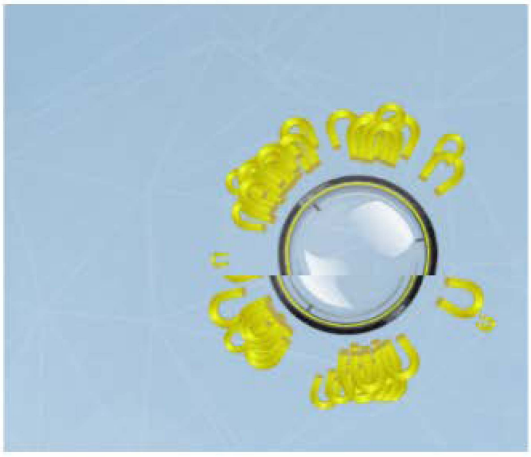

FIG 2. Swarm of Networked Users

This begs the question, can humans swarm? And if so, can we achieve similar benefits? The answer to both question appears to be yes. My personal experience with swarms has been at the Silicon Valley startup Unanimous A.I., which has been developing technologies that enable online human swarming. Their recently published studies have shown that swarming allows groups to make predictions and craft estimates that are more accurate than those achieved by polls, votes, surveys, and traditional forms of group decision making. For example, swarms of networked users have been shown to make accurate predictions about the outcome of sporting events, the price of commodities, and the winners of awards like the Oscars. But amplified intelligence is only one reason why Mother Nature evolved the process of Swarm Intelligence. The other reason is enabling groups to converge on decisions that support the common good.

So, can we humans use swarming to make decisions that better reflect our common values? And more importantly, can swarming help human groups avoid the Tragedy of the Commons pitfall? Early testing suggests the answer is yes. A recent study by Unanimous A.I. compared the decisions made by networked groups, first as disconnected individuals, then as a unified swarm. The study was modeled as a traditional "Tragedy of the Commons" dilemma in which subjects are asked to choose a cash bonus, the awarding of which is dependent upon the behavior of the full group. The test engaged 18 randomly selected online users, each paid $1.00 for their participation. All were told they would get an added bonus of $0.30 or $0.90 – they simply had to indicate on a blind survey which bonus they wanted. Of course there was a catch: if more than 25% of the group asked for $0.90, then nobody would get anything. This means oversubscription of the $0.90 option would defeat their common interests. And that's exactly what happened – a whopping 67% of the group asked for a $0.90 bonus on the survey, well beyond the 25% threshold. Thus, nobody received a bonus, the group failing to achieve their common interest.

The test was repeated by forming a real-time swarm, rather than taking a survey. This was done using an online platform called UNU. The interface allows networked groups to answer questions as a real-time *human swarm*, collectively exploring a decision-space and converging on a preferred solution.

Because people can't *waggle dance* like bees, the UNU platform, was designed to provide a humanfriendly interface that enables the same type of synchronous feedback loops. It works by allowing networked users to collaboratively move a graphical puck to select an answer, each person controlling their own small magnet to influence the direction and speed of the overall system. With everyone pulling in real-time, adapting hundreds of times per minute, a unified system emerges that reflects the collective will of the swarm. In this experiment, any of the users pulling towards $0.30 at the end of the decision would get that bonus, while anyone pulling towards $0.90 get that bonus. Thus, users were able to pursue their individual interests while helping to guide the overall swarm.

FIG 3. Snapshot of a Swarm in Action us

The results were inspiring: The swarm configured itself such that 24% of the total pull on the puck was towards $0.90, with 70% of the total pull towards $0.30, and 6% abstaining. Figure 3 shows a snapshot of that swarm in action, each magnet controlled by a unique user as the group worked together to move the puck. It's important to stress that the users could only see the puck and their

own magnet, but not the full swarm of other magnets. This means they had no direct indication of how many users were pulling in each direction. Still, the group, when functioning as a unified system, connected by real-time feedback loops, avoided the Tragedy of the Commons pitfall, instead converging on a solution that was best for group as a whole.

What does Swarm Intelligence mean for our future? It could point us to new methods for reaching group decisions – methods that encourage groups to combine their individual knowledge, opinions, and interests in support of the common good, avoiding entrenchment and stagnation. Further, human swarming could enable groups to reach decisions that better reflect our core morals and values, even in large groups where collective moralities often falter. And while many people still refer to the "hive mind" in a pejorative sense, I am now convinced this stems from misconceptions about how natural swarms work. The fact is, swarming is Mother Nature's brand of democracy, resulting from millions of years of evolution, and driven by a single selective motivator – to enable groups to work together for the good of the population as a whole.

Looking further out, online human swarms may be a path to super-intelligent systems. After all, a single honeybee lacks the intellectual capacity to even consider a complex problem like selecting a new home site for a colony, and yet swarms of bees have been shown to not only solve that multi-faceted problem, but find optimal solutions. If we humans could form similar swarms, we may be able to achieve similar boosts in intellect, solving problems that we currently, as individuals, can't even conceive. This is not only an exciting way to build smart systems, it's a path that keeps humans in the loop – ensuring that any super-intelligence that emerges has our core values and interests at its core.

SOURCES:

Cohen, Tara R., et al. "Group Morality and Intergroup Relations". Cross-

Cultural and Experimental Evidence Personality and Social Psychology Bulletin, Vol. 32, No. 11. (Nov 2006), pp. 1559-1572

Seeley, Thomas D., et al. "Stop signals provide cross inhibition in collective decision-making by honeybee swarms." Science 335.6064 (2012): 108-111.

K.m. Passino, T.F. Seeley, P.K. Visscher, Swarm Cognition in honeybees, Behav. Ecol. Sociobiol. 62, 401 (2008).

J.A.R. Marchall, R. Bogacz, A. Dornhaus, R. Planque, T.Kovacs, N.R. Franks. On optimal decision making in brains and social insect colonies,

J.R. Soc Interface 6,1065 (2009).

I.D. Couzin, Collective Cognition in Animal Groups, Trends Cogn. Sci. 13,36 (2008).

Seeley, Thomas D., Visscher, P. Kirk. Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making. Behavioral Ecology and Sociobiology 54 (5) 511-520.

Rosenberg, L.B, "Human Swarms, a real-time paradigm for collective intelligence." Collective Intelligence 2015, Santa Clara CA.

Rosenberg, L.B., "Human Swarms and collective intelligence." – European Conference on Artificial Life 2015, pp. 658-659

Rosenberg, L.B., "Human Swarming, a real-time method for Parallel Distributed Intelligence." – IEEE, 2015 Swarm/Human Blended Intelligence, Cleveland OH, USA.

https://commons.wikimedia.org/wiki/File:Bee_Swarm.JPG

About the Author:

Louis Rosenberg received B.S., M.S. and Ph.D. degrees from Stanford University. His doctoral work produced the *Virtual Fixtures* platform for the U.S. Air Force, the first immersive Augmented Reality system created. Rosenberg then founded Immersion Corporation (Nasdaq: IMMR), a virtual reality company focused on advanced interfaces. More recently, he founded